SCIENTIFIC & COMPUTATIONAL CHALLENGES OF THE FUSION SIMULATION PROJECT (FSP)

SciDAC 2008 CONFERENCE

W. M. TANG
Princeton University

Seattle, Washington

13-17 July 2008

ITER Goal: Demonstration of the Scientific and Technological Feasibility of Fusion Power

Further R&D is Needed to Make Fusion Practical

- **ITER is a truly dramatic step. For the first time the fusion fuel will be sustained at high temperature by the fusion reactions themselves.**
  - **Today:** ***10 MW(th) for 1 second with gain ~1***
  - **ITER:** ***500 MW(th) for >400 seconds with gain >10***

- **Many of the technologies used in ITER will be the same as those required in a power plant but additional R&D will be needed.**
  - **"DEMO":** ***2500 MW(th) continuous with gain >25, in a device of similar size and field as ITER***
    - ⇒ **Higher power density**
    - ⇒ **Efficient continuous operation**

- **Strong R&D programs are required to support ITER and leverage its results.**
  - **Experiments, theory, *computation*, and technology that support, supplement and benefit from ITER.**

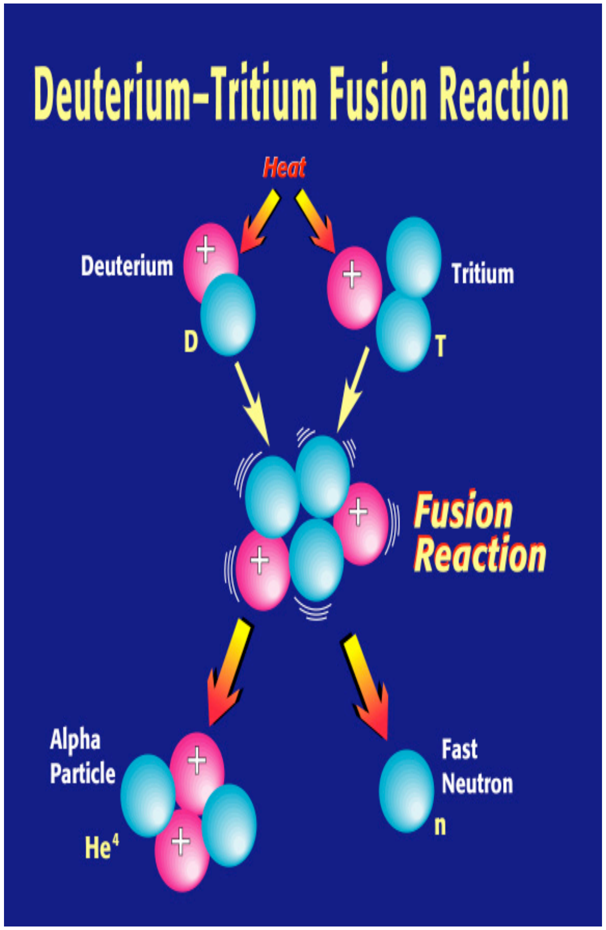

Fusion can be an Abundant, Safe, & Reliable Energy Source for the Future

- ***Exciting potential for long term availability of low cost fuel***
  - **No geopolitical instability due to competition for energy resources.**
- ***Reduced pollution and global climate change***
  - **No CO_2 production and no acid rain**
- ***Safe & reliable energy source***
  - **Short-lived radioactive waste and no possibility of runaway reactions/meltdown & low risk of nuclear proliferation**
- ***Estimated to be cost-competitive with coal & fission in future***

Complements nearer-term energy sources

Progress in Magnetic Fusion Research

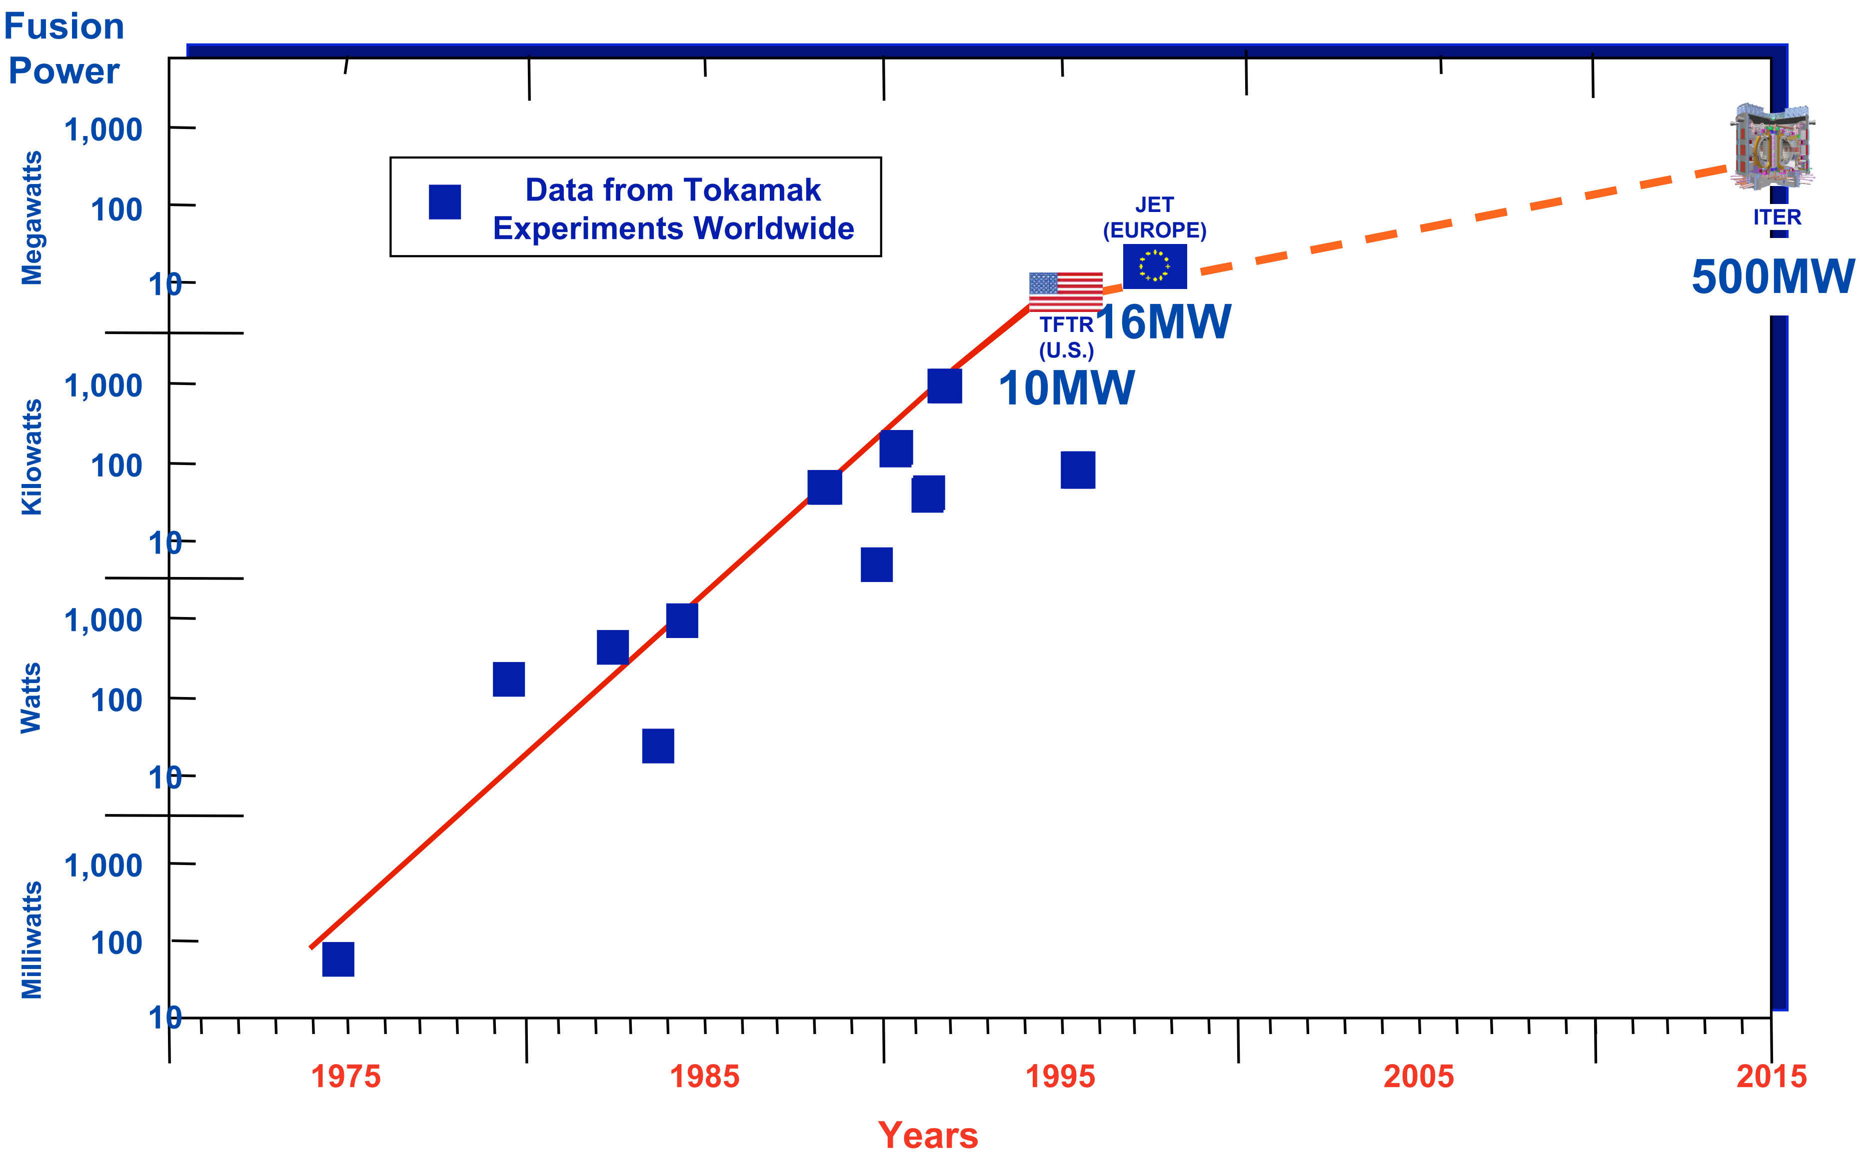

Situation Analysis: Why is FSP a Compelling & Timely Investment?

- <u>David Keyes</u>: "This is a ***historic opportunity for simulation***. The FSP arrives 'just in time' to help deliver a form of civilization's arguably most important technology: essentially inexhaustible, essentially proliferation-free carbon-neutral energy – *the summum bonum*."

- FSP lies in the crosshairs of priorities #1 and #2 of *Facilities for the Future of Science* and also of the 2007 *"E3"* initiative

- U.S. currently positioned to lead in large-scale integrated simulations of magnetically confined fusion plasmas

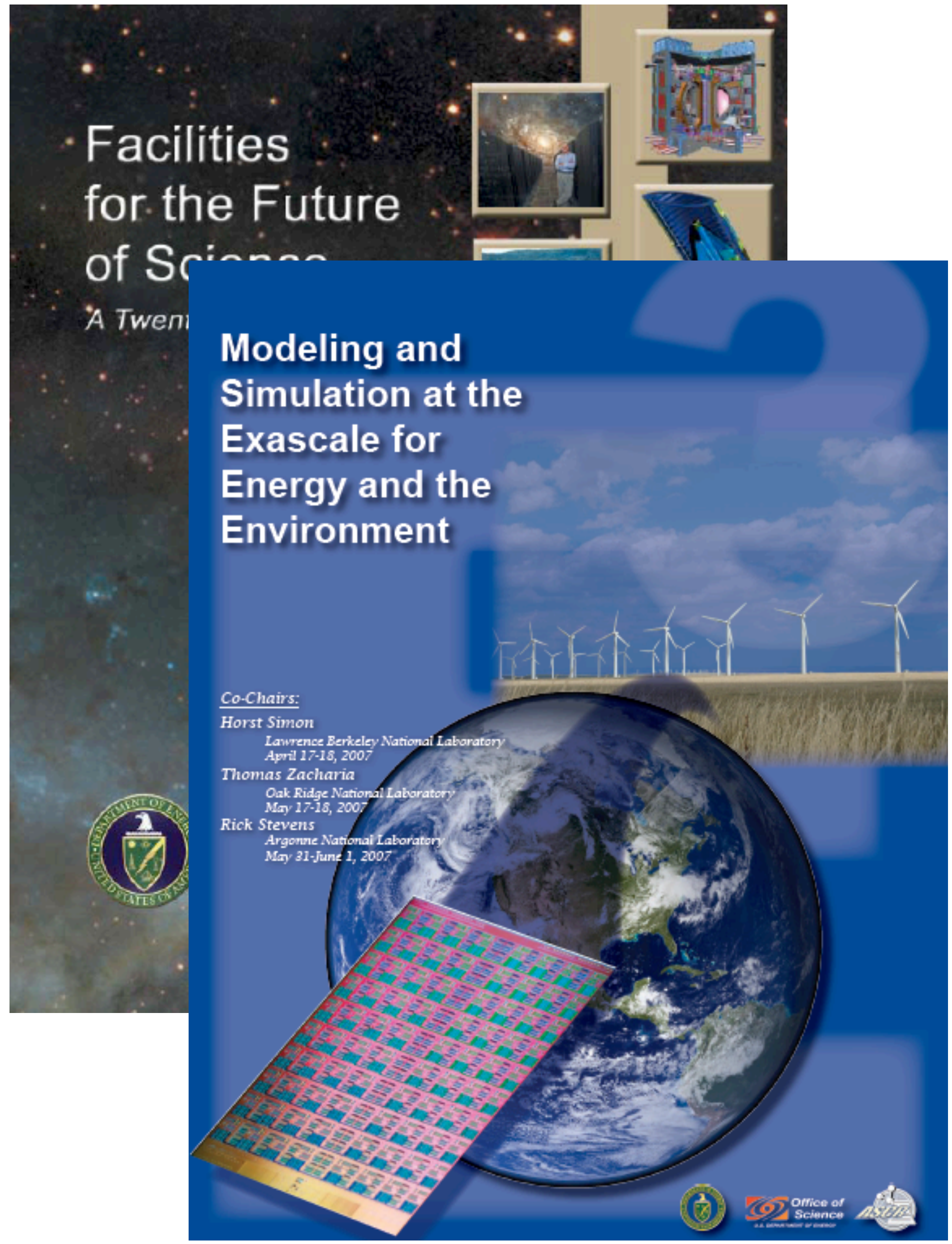

Unique Opportunity for US Leadership

- ***Critical need*** **for reliable predictive simulation capability for** ***ITER***
- **LCF's moving rapidly toward petascale & beyond** ***computing resources***
- **Interdisciplinary** ***collaborative experience*****, knowledge, & software assembled over 8 years under SciDAC plus OFES and OASCR base research programs**

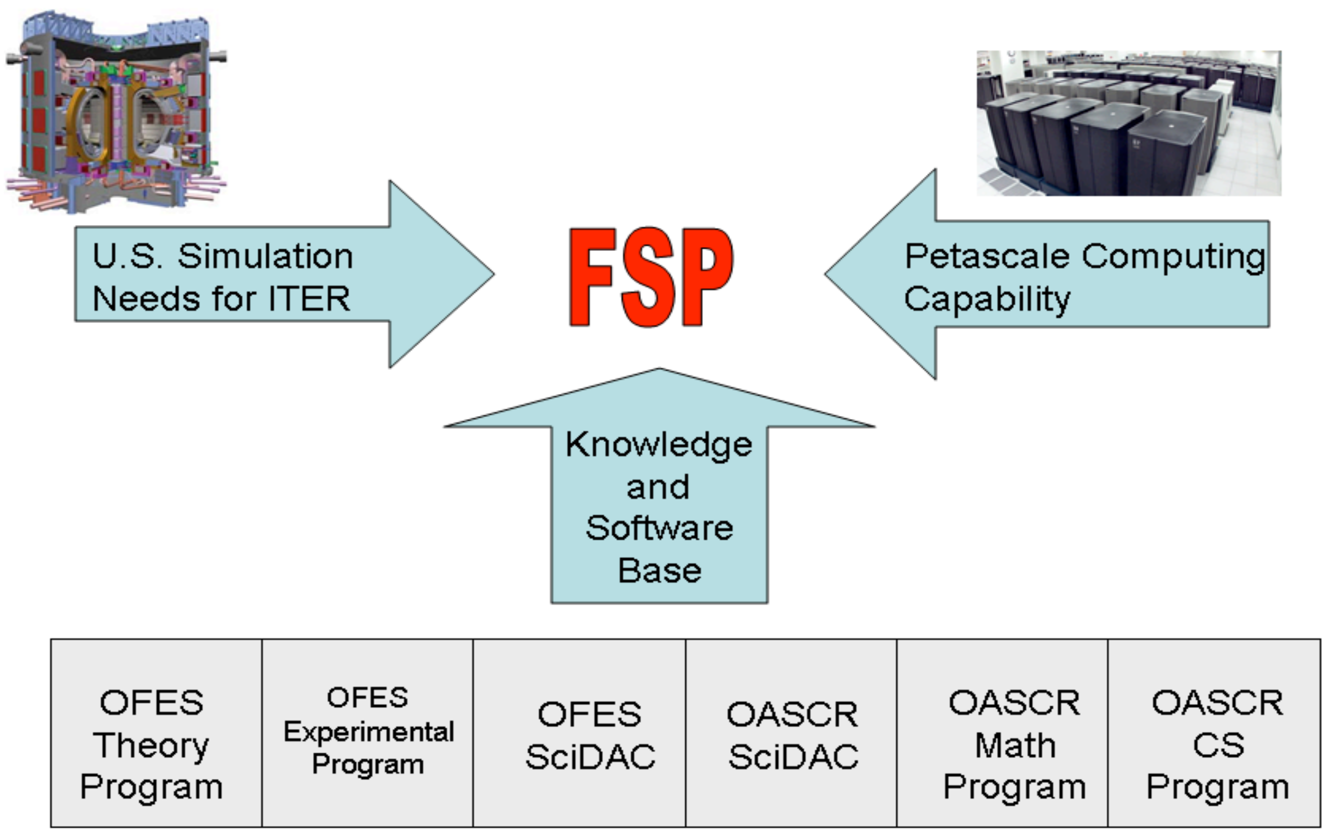

"Fusion Simulation Project" (FSP)

- <u>**Fusion Simulation Project**</u>

-targeting world leading US role in this area with impact on ITER & beyond

- FSP Dahlburg Report: 2002
- FSP Post Report: 2004
- FSP Workshop Report: July '07
- FESAC FSP Panel Report: October '07
- OFES has accepted FESAC FSP Panel recommendation for "Project Definition Phase" -- commencing in FY'09
- ASCAC FSP Panel Report due Aug. '08

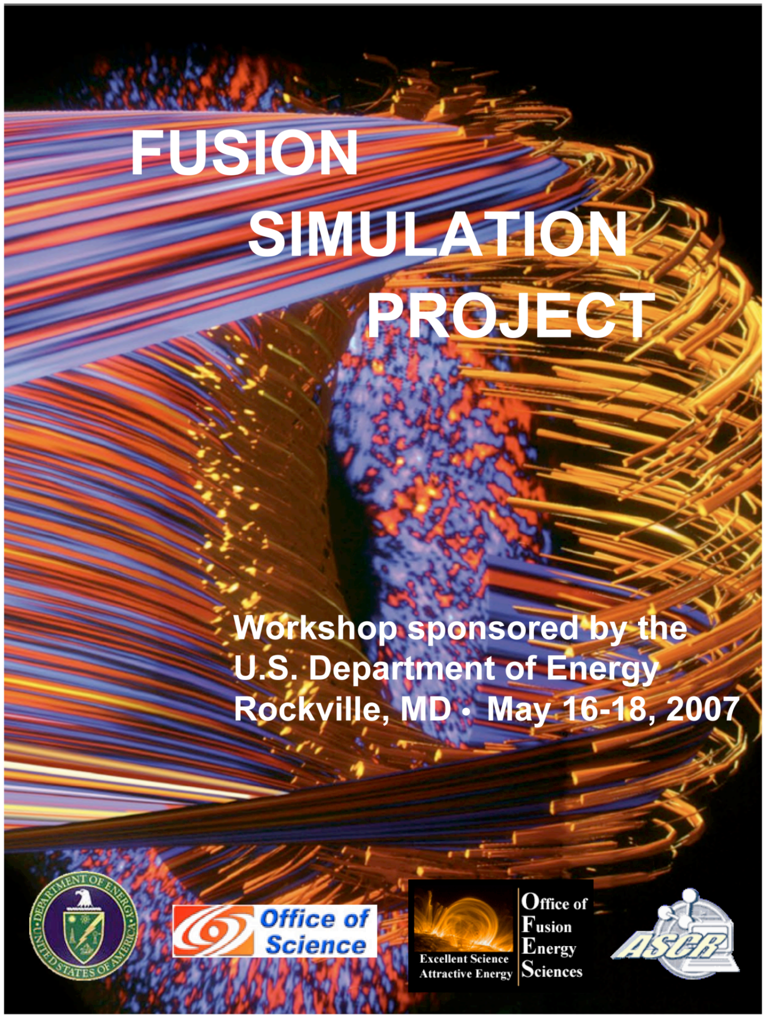

Fusion Simulation Project (FSP)

• **MISSION:** ***"Produce a world-leading predictive simulation capability that will be of major benefit to the science and mission goals of the US Fusion Energy Science Program."*** **(DOE Energy Undersecretary R. Orbach)**

• **GOAL: Develop high-performance software for an integrated modeling capability that would embody the theoretical and experimental understanding of confined thermonuclear plasmas such as ITER.**

- need to integrate into an effective framework the large number of codes and models that presently constitute separate disciplines within plasma science

• largest-scale codes addressing multi-scale physics of mostly individual phenomena in realistic 3D geometry

• integrated models with much smaller-scale lower dimensionality with significant empirical elements for interpretation & design of experiments

- need to effectively utilize petascale (& beyond) multi-core supercomputers with associated algorithmic advances

- need to reliably predict the behavior of plasma discharges in toroidal magnetic fusion devices on all relevant time and space scales in context of self-consistent simulations that are validated vs. experimental data

NATURE OF SCIENTIFIC & COMPUTATIONAL CHALLENGES FOR FSP

– **Critical scientific issues for fusion come from *"gaps analysis"* of capabilities from computational science that traditional *theory or experiment, by themselves, cannot readily deliver***

– **Critical computational issues come from *"gaps analysis"* of *capabilities missing from current state-of-art tools* for dealing with critical scientific issues**

Magnetically-confined Plasmas in a Tokamak

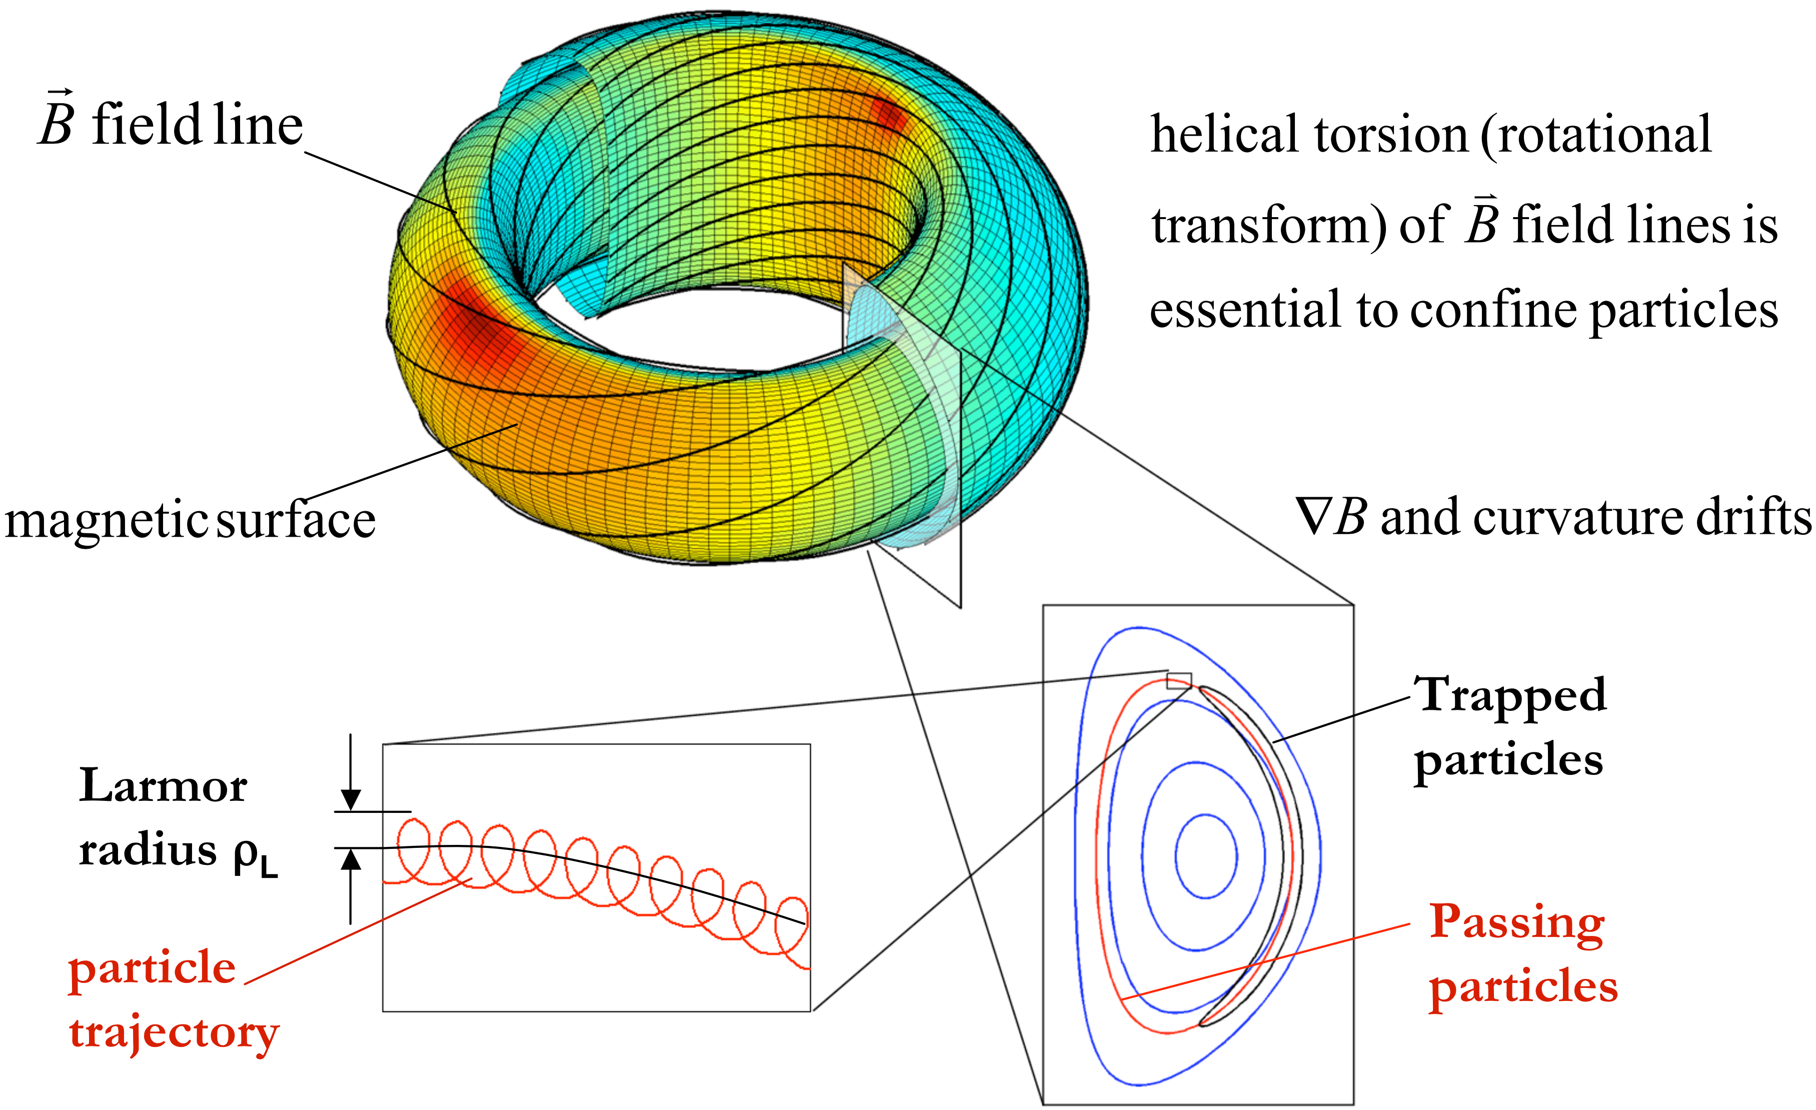

Elements of an Integrated Tokamak Modeling Code

- **Sawtooth Region (q < 1)**
- **Core Confinement Region**
- **Magnetic Islands**
- **Edge Pedestal Region**
- **Scrape-off Layer**
- **Vacuum/Wall/ Conductors/Antenna**

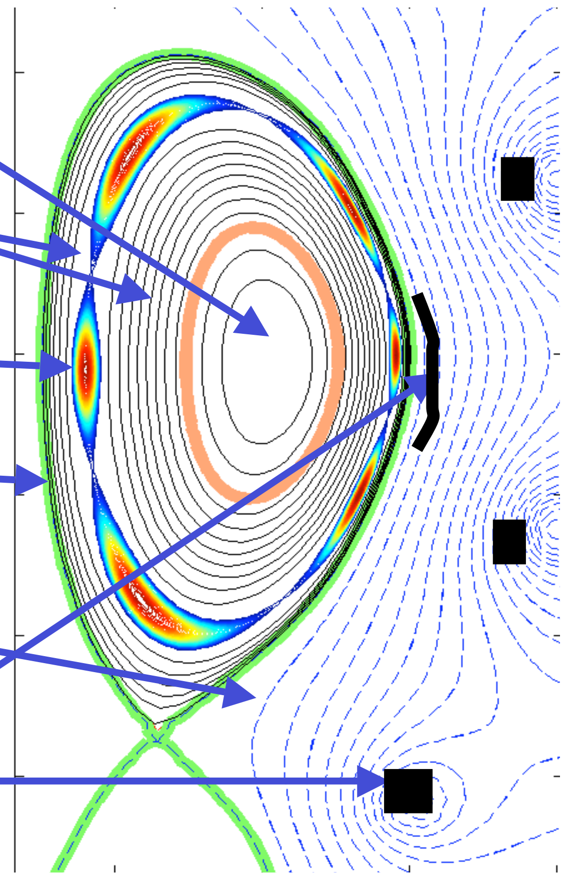

Core & Edge Transport

Plasma Turbulence

Large Scale Instabilities

MHD Equilibrium

Heating & Current Drive

Plasma-Wall Interactions

Atomic Physics

Radiative Transport

Energetic Particles

Huge Range of Spatial & Temporal Scales Present Major Challenge to Theory & Simulations

- **Overlap in scales often means strong (simplified) ordering not possible**

- ***Simulation at the <u>Petascale and beyond</u> will be essential for needed progress***

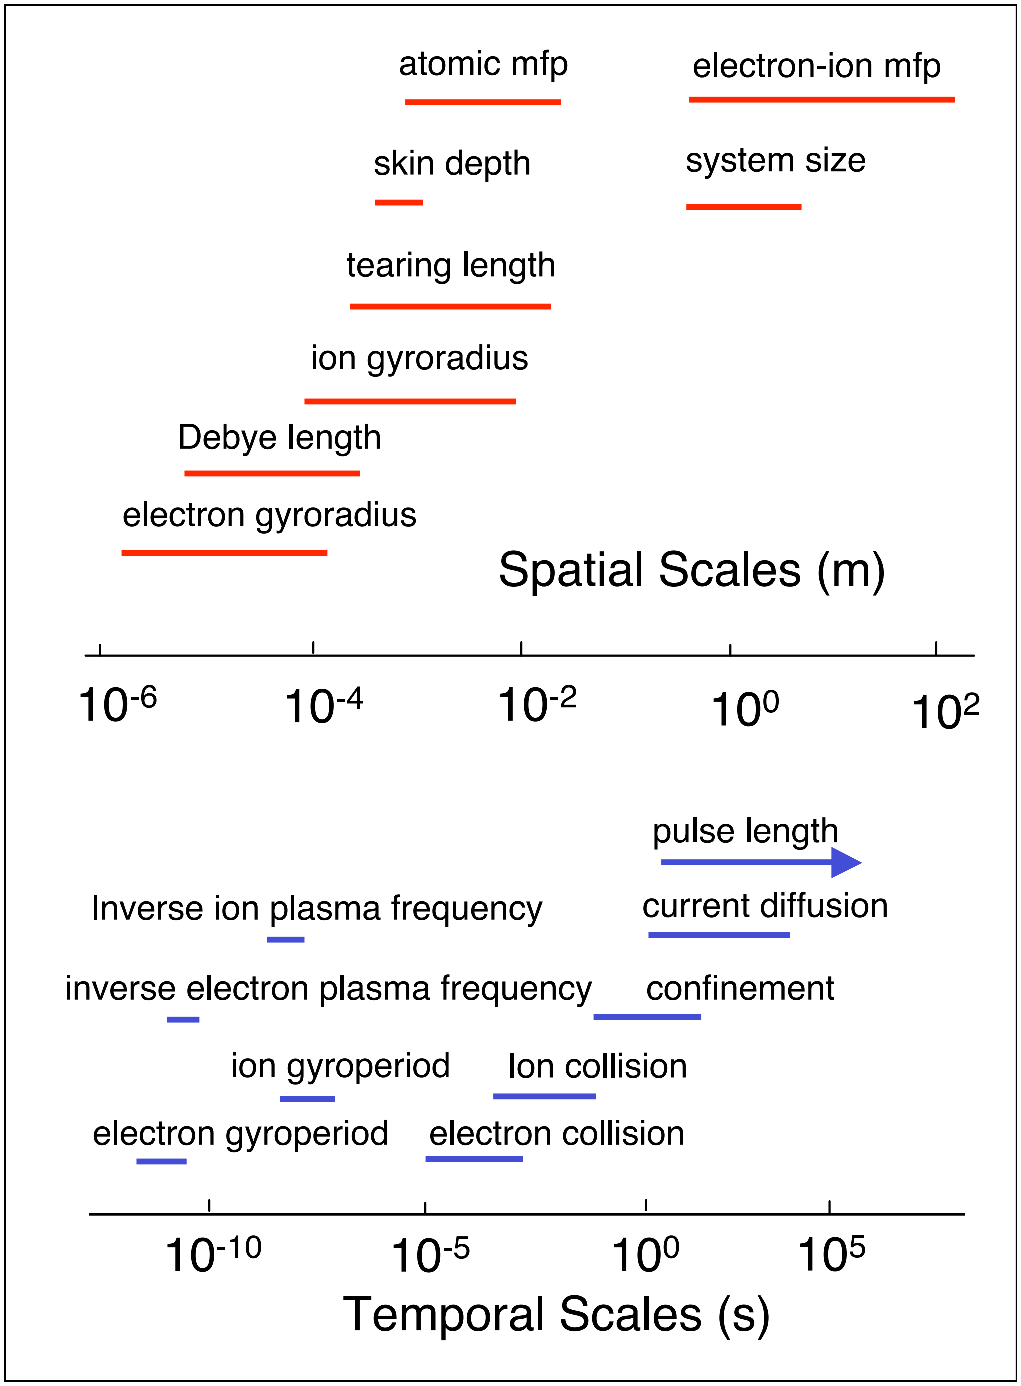

EXAMPLES OF IMPORTANT ADVANCES NEEDED FOR THE FSP

1. **Effective coupling of state-of-art codes for the *plasma core and edge* regions (e.g., C. S. Chang, et al. -- "Toward a first Principles Integration Simulation of Tokamak Edge Plasmas")**

2. **Effective coupling of state-of-art codes for *MHD dynamics and auxiliary heating of the plasma via RF waves* (e.g., D. Batchelor, et al. -- "Simulation of Wave Interactions with Magnetohydrodynamics")**

3. **Development of *advanced frameworks and workflow management* methods needed for code coupling (e.g., J. Cary, et al. -- "First Results from Core-Edge Parallel Composition in FACETS Project")**

4. **Development of more *realistic reduced models* based on results obtained from the direct numerical simulation (DNS) type major codes which use petascale capabilities**

5. **Development of appropriate *verification and validation* (V&V) effort to ensure reliable predictive capability**

KEY SCIENTIFIC CHALLENGES FOR THE FSP

(most urgent issues for burning plasmas and ITER operation)

1. **Disruption Effects & Mitigation** - ***Large-scale macroscopic events producing rapid termination of plasma discharges***
   - **Need to avoid since ITER can sustain only a limited number of full-current disruptions**
   - **Need to predict the onset of a disruption and to mitigate/minimize associated damage if it occurs**
2. **Pedestal (steep-spatial gradient) Formation and Transient Heat Loads on Plasma Periphery (divertor region)**
   - **Need to predict onset and growth of pedestal since its height is observed to control confinement**
   - **Need to predict frequency and size of Edge Localized Modes (ELMs) crashes to mitigate damage to the divertor and to plasma facing components**
3. **Tritium Migration and Impurity Transport**
   - **Need to predict tritium behavior since it can be hard to remove**
   - **Need to predict impurity influx and transport since they can dilute D-T fuel and degrade fusion power production**

KEY SCIENTIFIC CHALLENGES FOR THE FSP

(most urgent issues for burning plasmas and ITER operation)

4. Performance Optimization & Scenario Modeling

– **Need to optimize performance (including sustaining maximum fusion power production) while planning experiments since each ITER discharge will cost about $1M**

– **Need to control plasma current and pressure in more challenging scenarios -- moving from present experiments (10's of seconds duration) to ITER discharges dominated by alpha-self-heating and lasting thousands of seconds**

5. Plasma Feedback Control *- Burning plasma regime is fundamentally new with stronger self-coupling and weaker external control*

– **Need to design real-time feedback control to avoid disruptions and to optimize the performance of burning plasma experiments near operational limits**

– **Need to control edge localized modes (ELMs) since they can damage the divertor and impact the rapid erosion of plasma facing components**

NOTE: *Items (1) thru (3) focus on improved scientific understanding of physical processes [demanding integration of a few "1st principles solvers" with high physics fidelity] while (4) and (5) focus on new tools for operational control [requiring integration of a large number of reduced dimensionality models].*

Plasma Feedback: ELM Control/ Mitigation

Amplitude of uncontrolled ELM heat pulse in ITER expected to be order of magnitude above tolerable level for divertor plasma facing components

Magnetic Control (DIII-D tokamak)

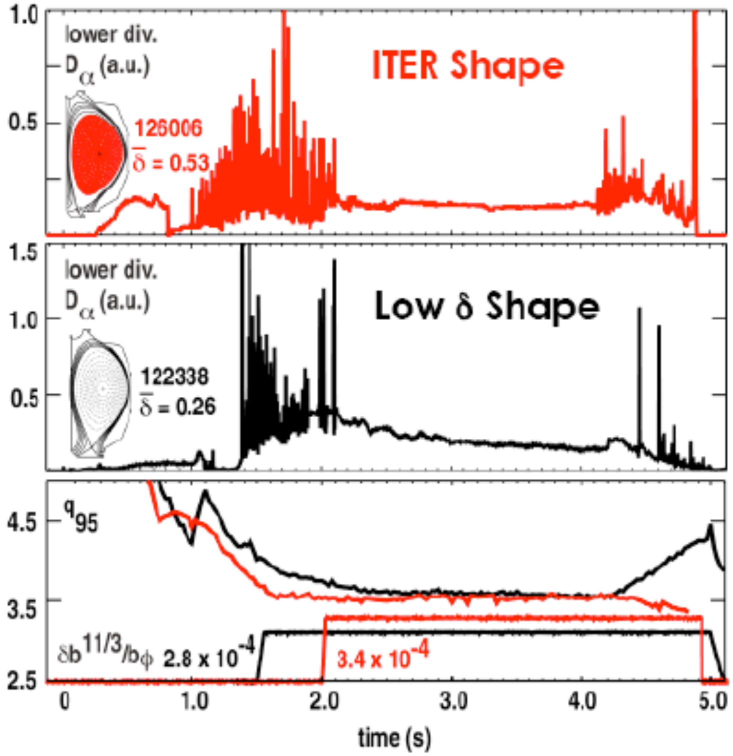

Pellet Pacemaking

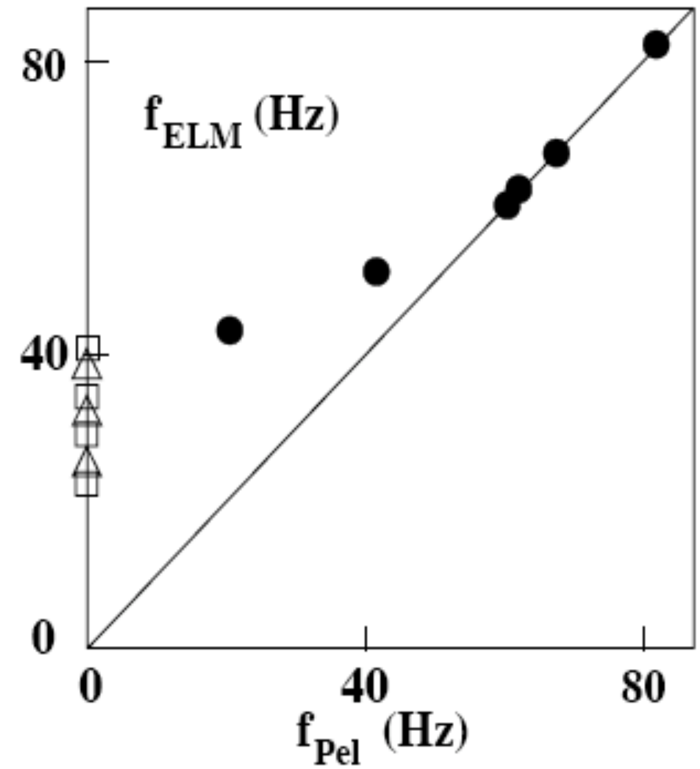

- **Two principal approaches currently under development for ITER:**
  - ***-- edge ergodization by Resonant Magnetic-Field Perturbation (RMP) coils***
  - ***-- pellet pacemaking***

Modelling of Pellet Fuelling Effectively Utilizes AMR *(R. Samtaney)*

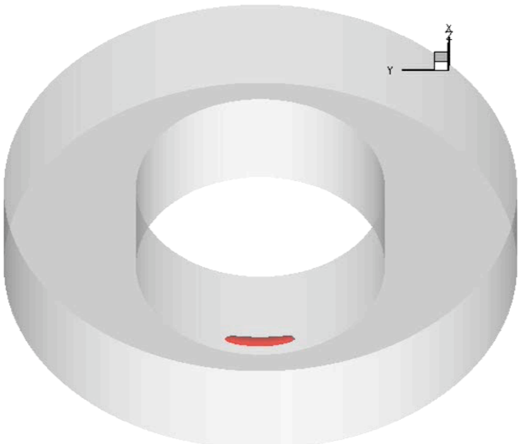

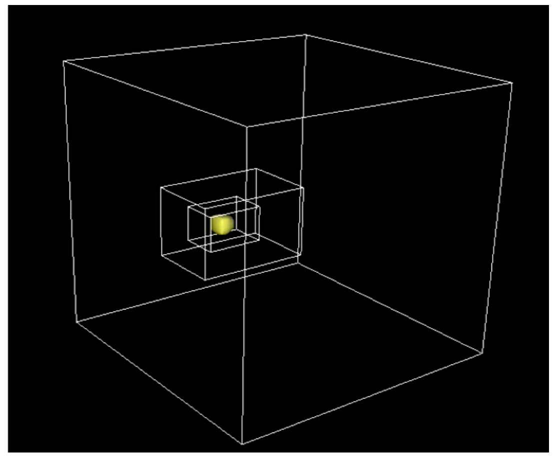

Adaptive Mesh Refinement

Inside Pellet Launch →

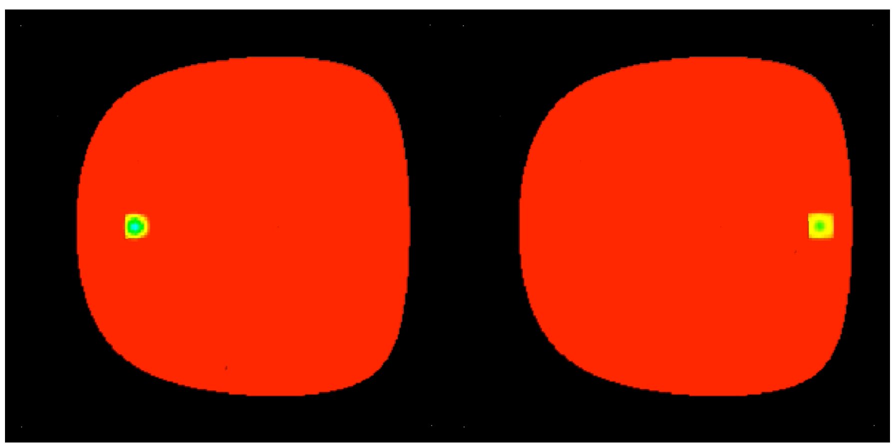

← **Outside Pellet Launch**

Challenges of Simulating Pellet-pacing of ELMs

- **Formidable multi-scale/multi-physics problem: resolving both pellet physics and ELMs**
  - **Small scales require adaptive mesh refinement (AMR) to resolve the pellet**
  - **Long time simulations (such as those for multiple pellet injections/ELM cycles) require development of implicit methods**
  - **Kinetic (long-mean-free-path) dynamics needed to properly model heat transport to the pellet-ablated cloud in complicated edge region -- (possible integration of appropriate kinetic models as “plug-ins”)**
- **Sophisticated numerical algorithms and software needed to address this computationally challenging integrated modeling problem**
  - ***Several SciDAC Centers collaborating to develop & enhance Chombo-based AMR MHD code with detailed pellet ablation physics coupled to ELMs with the objective of simulating pellet-induced ELM crashes***

Verification & Validation FSP Challenges

- **Establishing the physics fidelity of advanced physics modules demands proper Verification & Validation (V&V)**

• Verification assesses degree to which a code correctly implements the chosen physical model

--- more than "essentially a mathematical problem" --- Special emphasis should be placed on code verification via cross-code benchmarking and comparisons with theoretical predictions

•Validation assesses degree to which a code describes the real world, e.g.

-- *Development & Application of "Synthetic Diagnostics" in RF applications (Ref. -- Ref. P. Bonoli, et al. SciDAC'07 Conf. Proceedings)*

-- *C. Holland's presentation on "Validating Simulations of Core Turbulence Simulations: Current Status and Future Directions"*

Advanced Scientific Codes --- "*a measure of the state of understanding of natural and engineered systems*" (T. Dunning)

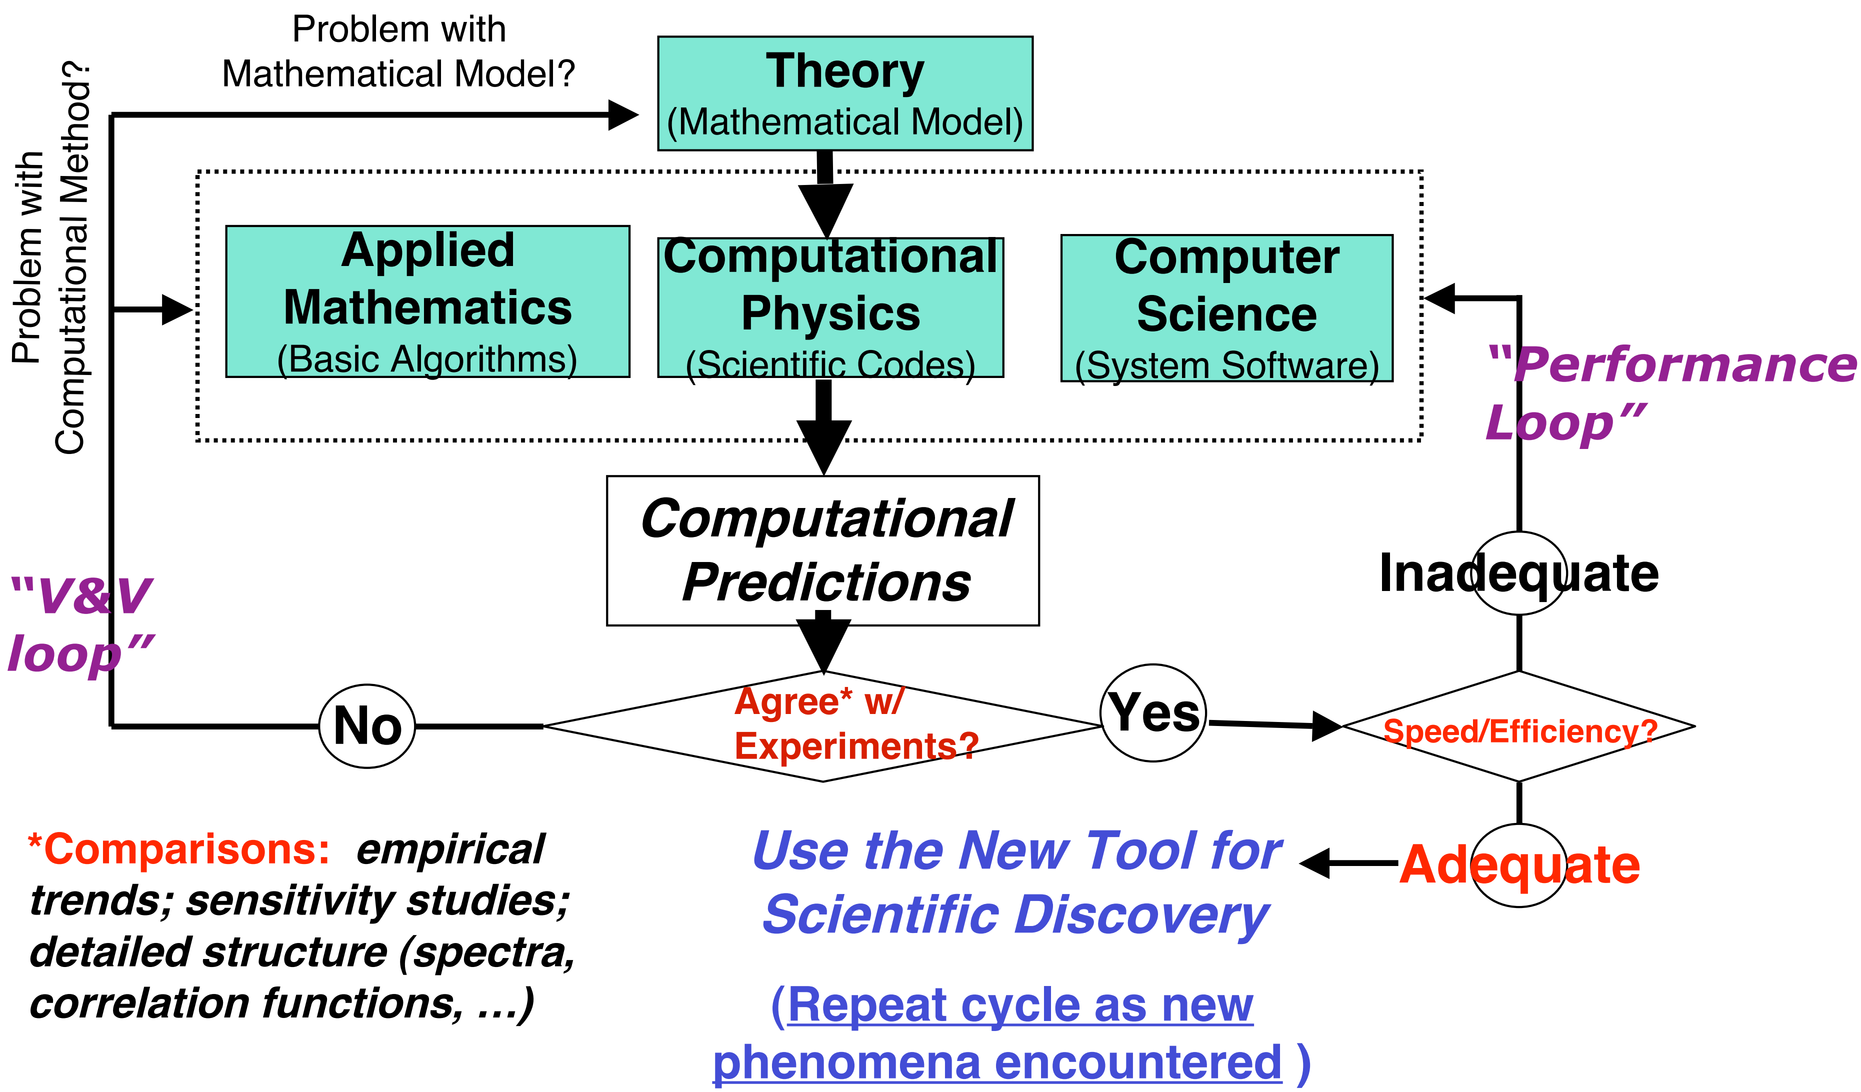

*Comparisons: *empirical trends; sensitivity studies; detailed structure (spectra, correlation functions, …)*

Mathematical and Computational Enabling Technologies Challenges for FSP

- **Risk quantification associated with each key part of HPC FSP software with appropriate identification of backup solutions and/or recovery methods**

- **Code flexibility to weather significant evolution of hardware architecture as well as that of associated systems software**

- **More efficient non-linearly scalable MHD codes to address challenges associated with anisotropy and stiffness**

- **Improved particle-in-cell and Eulerian (continuum) methods for addressing kinetic dynamics in complex geometry**

- **Gyrokinetic models with large gyroradius > Δx (with associated Poisson convolution issues)**

- **Data management, mining, advanced visualization, efficient storage capabilities for massive amounts of data**

Applied Mathematics Challenges for FSP

- Solution-adaptive mesh refinement (AMR) methods with higher order in space & time fitting complex geometry

- Nonlinearly implicit preconditioned Jacobian-free Newton-Krylov methods
  Ref: L. Chacon "Scalable Parallel Implicit Solvers for 3D Magnetohydrodynamics"
  - For individual research codes (e.g., extended MHD)
  - For coupling codes implicitly (multi-physics applications)

- Multi-scale methods with each phenomenon computed on appropriate scale with effective transfers to other scales

- Evolution to a million threads and beyond, pressuring algorithms to:
  - Communicate and synchronize less
  - Store less (and recompute more)
  - Copy data between different structures less

Modeling of RF Auxiliary Heating Needed for ITER Requires Petascale to Exascale Resources (P. Bonoli)

"Scientific Discovery" **- Full-wave terascale RF simulations (AORSA, TORIC) reveal significant spectral broadening due to diffraction –** ***not described by standard ray tracing techniques.***

Single antenna mode can be simulated on the CRAY XT4 JAGUAR at ORNL in 1 hour with 4096 processors.

Coupled Fokker Planck – Full wave simulation (CQL3D + TORIC) with full antenna spectrum requires about 600,000 CPU hours in a present day sized device.

Calculation size increases by a factor of 10^4 for ITER-sized device

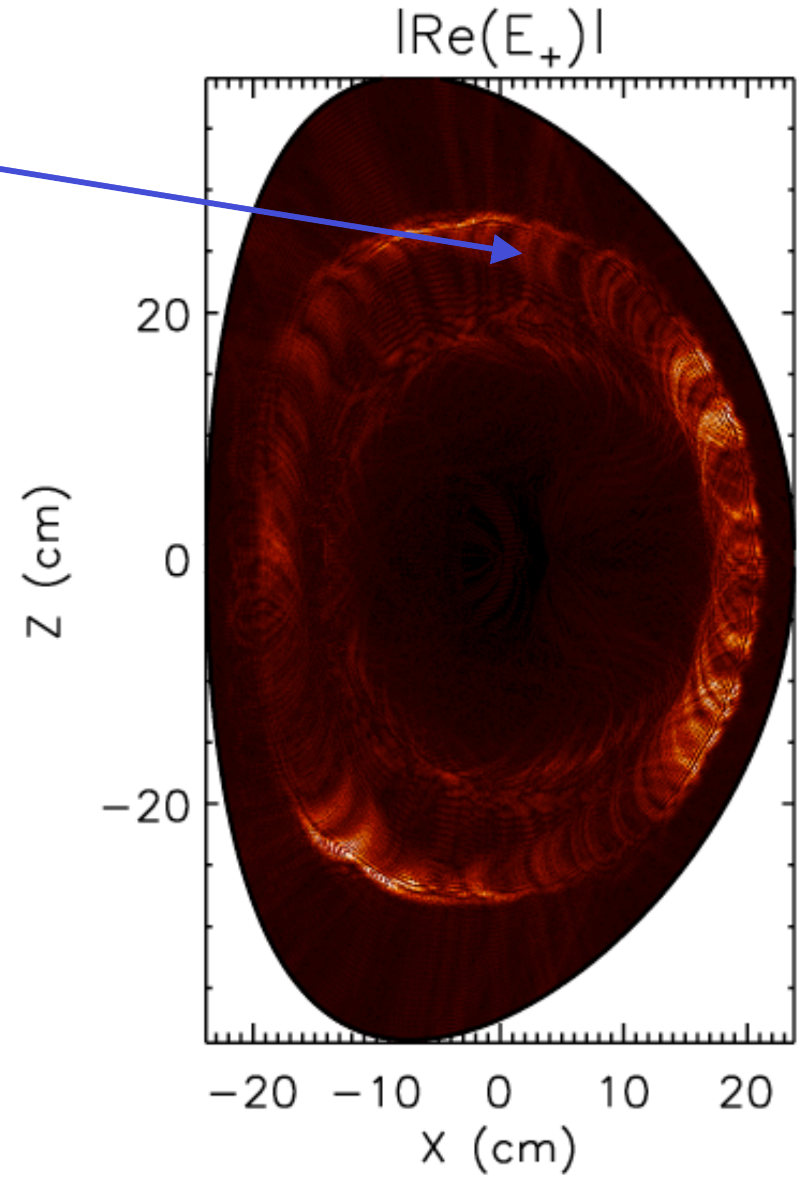

Microturbulence in Fusion Plasmas: *Size & Cost of reactor from balance between confinement & fusion self-heating rates*

- *"Scientific Discovery"* - Transition to favorable scaling of confinement for both *ions* and *electrons* now observed in simulations for ITER plasmas
- *Electron transport* less understood but more important in ITER since fusion products first heat the electrons
- Simulation of electron turbulence is more demanding due to shorter time scales and smaller spatial scales
- Recent GTC simulation of electron turbulence *used 28,000 cores for 42 hours in a dedicated run on Jaguar at ORNL producing 60 TB of data* currently being analyzed

Ref: Z. Lin, Y. Xiao, et al, Fusion poster session, Tuesday evening

Good news for ITER!

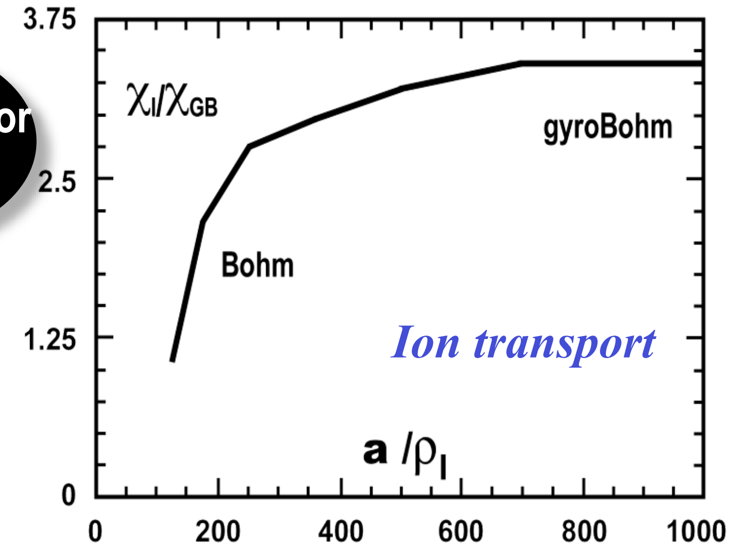

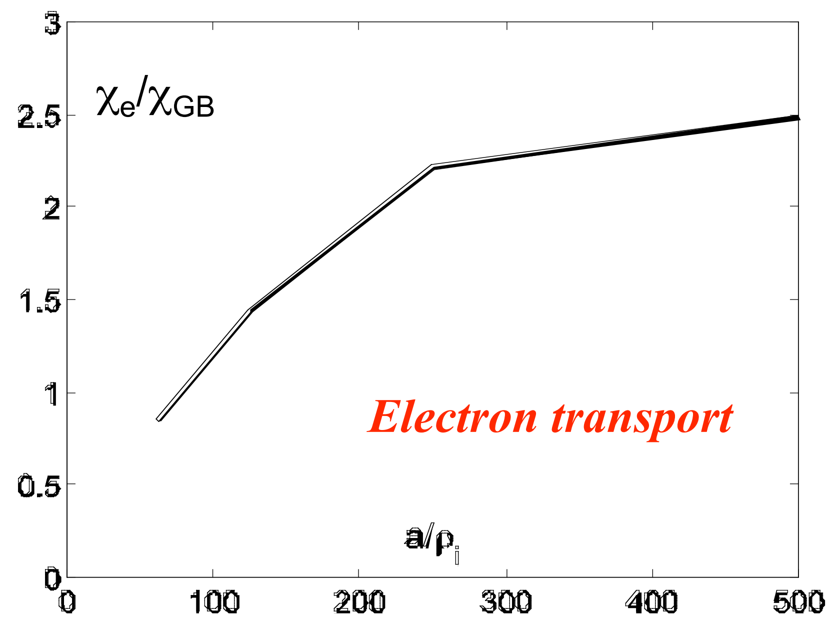

Recent High-Resolution Simulations

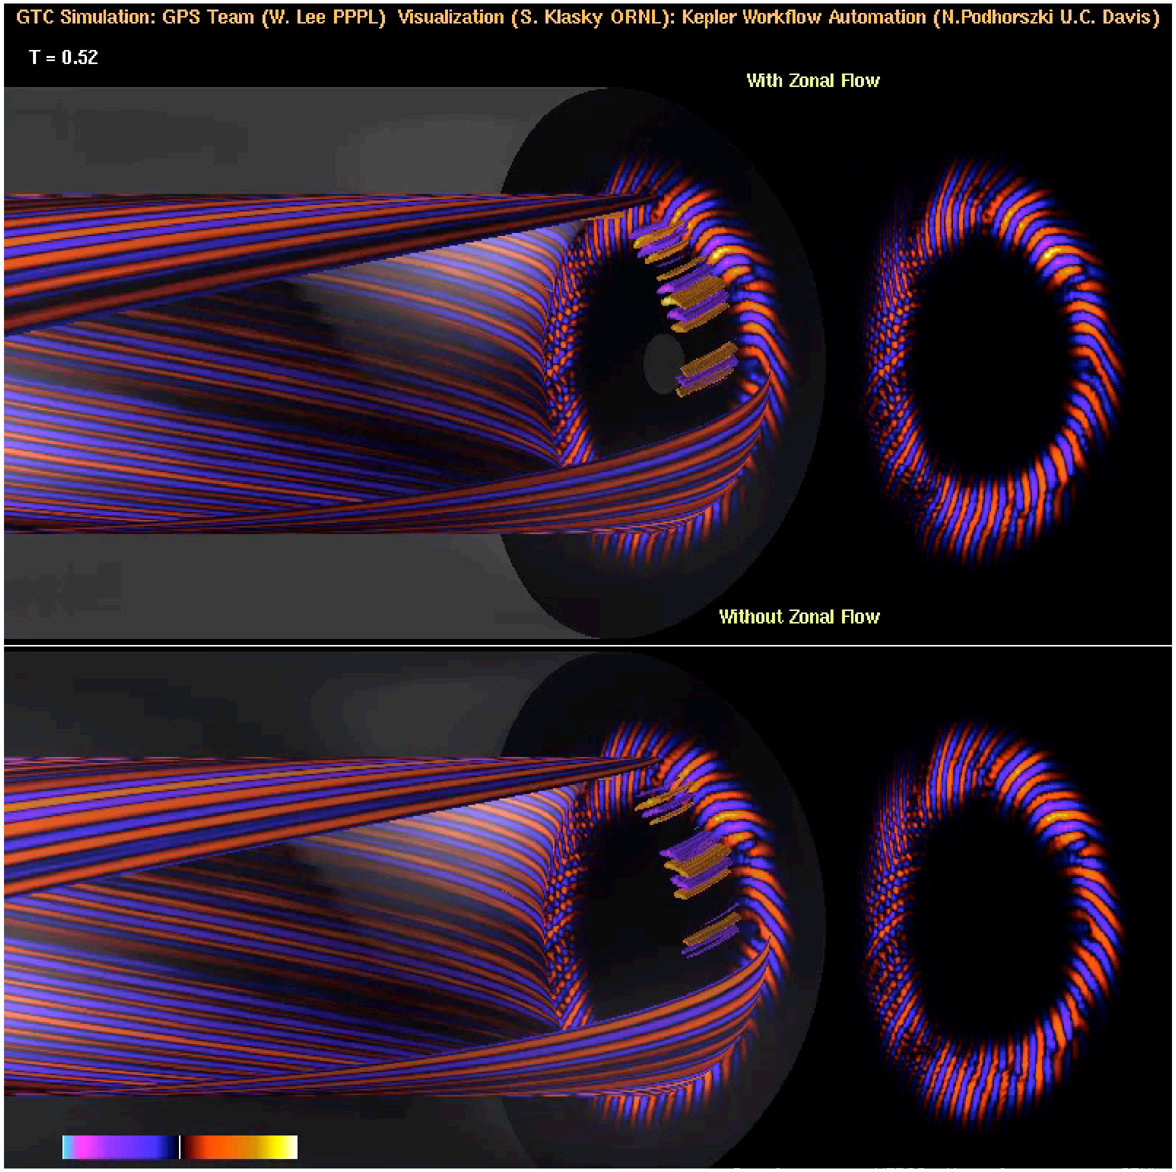

- High-resolution visualization from ***realistic shaped-cross section toroidal plasma simulations*** on leadership class computers
  - **[SciDAC GPS Center & ORNL's Jewel Milestone project (W. Wang, et al.)]**

- Efficiently generated via "Workflow Automation" -- ***automation of data movement, data reduction, data analysis, and data visualization*** **[SciDAC SDM Center's Kepler workflow project (S. Klasky, et al.)]**

EFFIS: End-to-end Framework for Fusion Integrated Simulation [SciDAC SDM (LBNL, ORNL, …..)]

- ***Enabling technology framework for HPC software on LCF's*** *including advanced data management, code coupling, I/O, visualization, & monitoring*

- ***Elements:***
  -- Workflow engine (e.g., Kepler)
  -- Adaptable I/O System (e.g., ADIOS)
  -- Provenance data-base support
  -- Wide-area data movement
  -- Code coupling
  -- Visualization
  -- Dashboard

- ***Application:*** *SciDAC CPES*

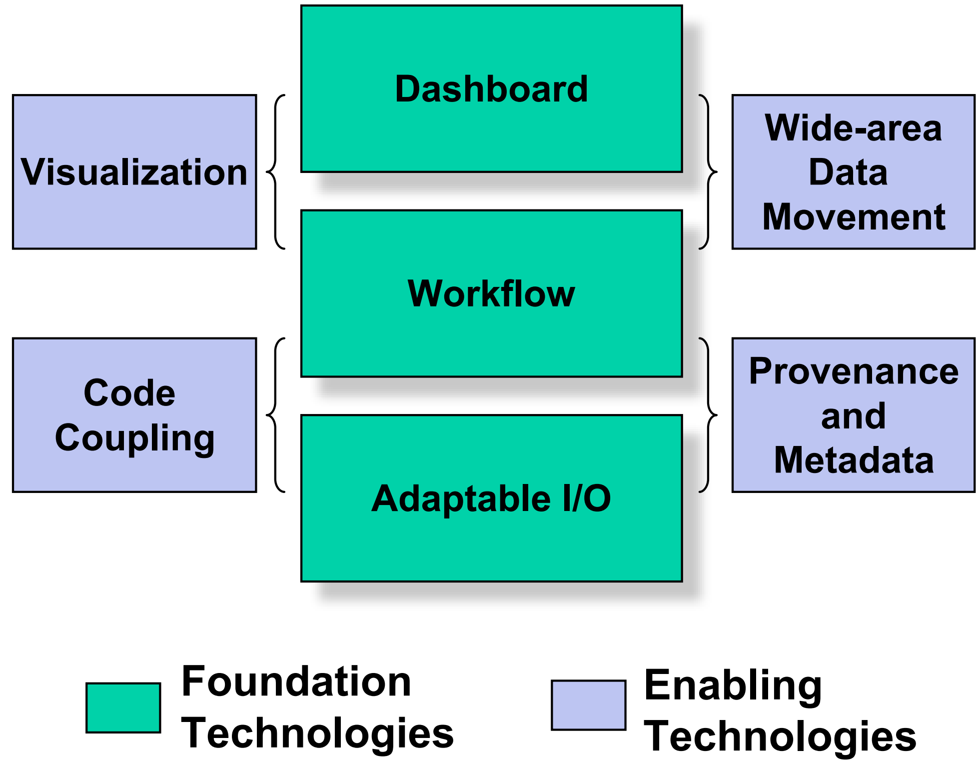

FSP: Lots of Exciting Physics & Algorithms Challenges!

Physics

• Disruptions: avoidance & mitigation
• Pedestal formation & transient divertor heat loads on plasma periphery
• Tritium migration & impurity transport
• Performance optimization & scenario modeling
• Plasma feedback control
• Verification & validation (including Synthetic diagnostics) to ensure reliable predictive capability
• Physics modeling components to address: (1) core and edge turbulence transport; (2) large scale (MHD) instabilities; (3) sources and sinks of heat, momentum, current and particles; and (4) energetic particle effects.

Algorithms

• Efficient multi-core algorithms -- including addressing evolution to million threads and beyond
• Multi-physics integration capabilities
• Advanced frameworks for code coupling
• Multi-scale computation of phenomena on appropriate scale with effective transfers to other scales
• Solution-adaptive mesh refinement (AMR) methods with higher order in space & time fitting complex geometry
• Nonlinearly implicit preconditioned Jacobian-free Newton-Krylov methods
• Management, analysis, advanced visualization, efficient storage capabilities for massive amounts of data

Necessary to attract, train, & assimilate best & brightest people

Concluding Comments

- **Progress in magnetic fusion research achieved has been dramatic -- leading to *ITER* ($10B burning plasma experiment)**
  - **located in France; supported by 7 nations representing over half of world's population**
  - **ITER targets 500 MW for 400 seconds with gain > 10 to demonstrate *technical feasibility of fusion energy***
  - ***US can play the lead role in using advanced computation to harvest knowledge from ITER***
- ***DEMO* (demonstration power plant) targets 2500 MW with gain of 25 -- demands *R & D with computation @ petascale and beyond as critical component***
- ***FSP* will target realistic simulations of fusion systems with unprecedented physics fidelity**

-- includes: (i) delivering shorter-term opportunistic HPC software tools (built largely from existing tools) & (ii) parallel longer-term development emphasizing new, more rigorous, more engineered performance capabilities

-- exciting advances for *predictive capabilities* will be driven by access to *LCF's* -- from terascale to petascale & beyond -- together with a vigorous *verification & validation program*